

AFRL-VA-WP-TP-2002-308

**UAV SWARMING? SO WHAT ARE
THOSE SWARMS, WHAT ARE THE
IMPLICATIONS, AND HOW DO WE
HANDLE THEM?**



Bruce T. Clough

APRIL 2002

Approved for public release; distribution is unlimited.

This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

20020912 087

**AIR VEHICLES DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542**

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) April 2002		2. REPORT TYPE Conference Paper Preprint		3. DATES COVERED (From - To) 03/01/01 – 03/01/02	
4. TITLE AND SUBTITLE UAV SWARMING? SO WHAT ARE THOSE SWARMS, WHAT ARE THE IMPLICATIONS, AND HOW DO WE HANDLE THEM?				5a. CONTRACT NUMBER IN-HOUSE	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 62201F	
6. AUTHOR(S) Bruce T. Clough				5d. PROJECT NUMBER 232F	
				5e. TASK NUMBER FE	
				5f. WORK UNIT NUMBER PA	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Control Systems Development and Application Branch (VACC) Control Sciences Division Air Vehicles Directorate Air Force Research Laboratory, Air Force Materiel Command Wright-Patterson Air Force Base, OH 45433-7542				8. PERFORMING ORGANIZATION REPORT NUMBER AFRL-VA-WP-TP-2002-308	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AIR VEHICLES DIRECTORATE AIR FORCE RESEARCH LABORATORY AIR FORCE MATERIEL COMMAND WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7542				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL/VACC	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-VA-WP-TP-2002-308	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES Paper for AUVSI Unmanned System Conference, 9-11 July 02 in Orlando, FL. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.					
14. ABSTRACT (Maximum 200 Words) Swarming. Imagine unmanned aerial vehicles blotting out the sky, obliterating anything that moves, and injecting terror into every observer's heart. Impossible? Maybe not. The aerospace research community is working hard at developing UAV control technology that requires as little human supervision as possible, and concepts using swarms are receiving serious attention. Swarming is not just a scare tactic, but also a viable control technology for multiple autonomous vehicles that system designers can use. Swarming itself is a type of emergent behavior, a behavior that isn't explicitly programmed, but results as the natural interaction of multiple entities. Swarms used correctly could be terribly effective, but used the wrong way could be as vulnerable as gnats to Raid™. This paper looks at what the author considers swarming to be, how you get it, what it might be (and might not be) good for, and the questions we need to answer on the way forward to evaluate and implement it.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 22	19a. NAME OF RESPONSIBLE PERSON (Monitor) Bruce T. Clough 19b. TELEPHONE NUMBER (Include Area Code) (937) 255-2831
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			

UAV Swarming? So What Are Those Swarms, What Are The Implications, And How Do We Handle Them?

Bruce Clough
Technical Area Leader: Control Automation
Air Force Research Laboratory
937.255.2831 *Bruce.Clough@wpafb.af.mil*

Abstract

Swarming. Imagine unmanned aerial vehicles blotting out the sky, obliterating anything that moves, and injecting terror into every observer's heart. Impossible? Maybe not. The aerospace research community is working hard at developing UAV control technology that requires as little human supervision as possible, and concepts using "Swarms" are receiving serious attention. Swarming is not just a scare tactic, but also a viable control technology for multiple autonomous vehicles that system designers can use. "Swarming" itself is a type of emergent behavior, a behavior that isn't explicitly programmed, but results as the natural interaction of multiple entities. Swarms used correctly could be as terribly effective, but used the wrong way could be as vulnerable as gnats to Raid_{TM}. This paper looks at what the author considers swarming to be, how you get it, what it might be (and might not be) good for, and the questions we need to answer on the way forward to evaluate and implement it.

Introduction

Imagine swarms, sheer terror. Imagine you're Helen running from the birds in *The Birds* [1]. Imagine you're an Army soldier faced with gigantic locusts in *Them* [2], imagine you're the poor cow wandering too close to the killer bee hive, or imagine you're the author who just put a lit firecracker in a hornet's nest [Note 1]. Now imagine killer UAVs swarming in for the kill. Far fetched?

Maybe not!

Swarming is one of the current "buzzwords" of choice in Department of Defense circles. We want mines to swarm. We want ground sensors to swarm, and we want aerial vehicles to swarm. We want to leverage nature to bring distributed control to our weapon systems. Bad idea? No. As with some buzzwords "swarming" is being a bit over-used today; however, the bio-inspired distributed control technology has good merits for some tasks, won't work well for others. Rather than being a general panacea or "snake-oil", swarming has its place in the repertoire of methods available to system designers.

This paper explores exactly what that place is, first defining swarms from a autonomous control system view, then comparing them with the other popular form of multiple autonomous system interaction – teams. In the end we explore six general tasks that swarming (at least for UAVs) seems to be effective for. Finally we lay out the challenges facing autonomous system designers who want to use swarms, challenges that define our future technology development focus. Swarms used correctly could be as terribly effective as the killer cockroaches in *Damnation Alley*[3], but used the wrong way could be as vulnerable as gnats to Raid_{TM}.

This paper is declared a work of the U.S. Government and not subject to copyright protection in the United States.

Background

If one hasn't noticed, there is a huge push within the Department of Defense to reduce the number of Airmen, Soldiers, Sailors and Marines exposed to hostile action. One of the ways we are doing this is by building and fielding unmanned weapons, weapons that are smart enough to do a task formerly done by a human. Progresses in mechanical miniaturization, cheap computing, and artificial intelligence have enabled the push.

One of the more active areas of controls research is the application of biologically-inspired control research to human projects, including unmanned vehicles [4]. The reasoning is simple – autonomous systems have existed for millions of years in the form of animals and plants seen all around us. We might learn a thing or two about control by watching those systems' behaviors and emulating them. The study of social insect behavior is an especially hot item – how do bees and ants, with their simplistic neural systems, interact to form complex colonies [7].

Couple the study of animal behavior to the rise of complexity theory – that very complex and elaborate structures emerge naturally from the interaction of simple rules and individuals [10] – and one can develop a methodology for developing autonomous system controllers. As will be shown later, a “swarm” is nothing more than one type of emergent behavior of simplistic individuals.

Sometime in the last few years someone has put two and two together and developed the concept of a group of simple, inexpensive unmanned systems doing the job of a highly trained human. This is the swarm of unmanned aerial vehicles (UAVs) that will blot out the sky and bring terror to men's hearts, or at least that's what's hoped!

So What Is Swarming?

Let's start with definitions. Here “Swarming” is not about numbers, but about the interactions between individuals in a group. If pure numbers were all it took then we could call the VII Corps in Desert Storm a swarm, a term that Gen Franks might disagree with. Swarming is a particular emergent behavior of simple autonomous individuals.

In a Rand report written about swarming [5], swarming is looked at in two ways: the group interaction amongst individuals, and swarming tactics that could be done by military groups (originally dispersed, the groups come together (or swarm) to accomplish an objective, then disperse again – somewhat akin to guerilla war tactics). The report then goes on to describe a system that uses decentralized control to command a swarm of low cost precision munitions. This paper concerns their first definition of a swarm, decentralized control of low cost autonomous systems. Readers wanting further information on military swarming tactics are encouraged to read Rand's report *Swarming & The Future of Conflict* [15].

So what would one consider a formal definition of swarming?

Formal Definition

Webster's Dictionary [6] gives this definition for “swarms”:

1a (1): a great number of honeybees emigrating together from a hive in company with a queen to start a new colony elsewhere (2) : a colony of honeybees settled in a hive **b** : an aggregation of free-floating or free-swimming unicellular organisms -- usually used of zoospores

2a: a large number of animate or inanimate things massed together and usually in motion : <swarms of sightseers> <a swarm of locusts> <a swarm of meteors> **b** : a number of similar geological features or phenomena close together in space or time <a swarm of dikes> <an earthquake swarm>

3: to form and depart from a hive in a swarm

4a: to move or assemble in a crowd, Tb : to hover about in the manner of a bee in a swarmM

When folks talk of UAVs that “swarm”. I believe that they are talking about 2a above, but they are certainly thinking about 2b!

Swarms - My definition

I have a less formal definition:

Swarming: *Getting a bunch of small cheap dumb things to do the same job as an expensive smart thing.*

Maybe a bit simplistic, perhaps blatant, but this is exactly what we all want to do. We want to get a bunch of \$1 robots to accomplish the same job as our current \$100M system. Not that there is anything bad about this! Of course, being technical types we’ve attached a more formal definition to our meaning of swarming:

Swarming: *A collection of autonomous individuals relying on local sensing and reactive behaviors interacting such that a global behavior emerges from the interactions.*

Note I said nothing about flying. Our usual thought about swarms and swarming is usually some type of flying social insect with a stinger attached. We are generalizing this to cover the emergent behaviors of any group of interacting individuals. Emergent behaviors? Yes, swarming, in our definition, is an emergent behavior. What’s that? Read on.

What Are The Technological Requirements For Swarming?

So what do we need to build a swarm? What’s the minimal set of technologies to bring together to enable swarms? We assert that all one needs is an appropriate set of individual reactive behaviors supported by local sensing that will interact to develop the group behavior of interest, and a simple architecture to manage those behaviors in the individual. In our views, just four:

- Emergent behavior of the interacting individuals in the swarm
- Simple reactive behaviors in the individuals
- Behavior architecture in the individuals allowing switches between reactive behaviors
- Local sensing with individual variability

Let’s review a bit on each:

Emergent Behaviors

The best way to describe emergent behaviors is by starting with a few examples. Flocking birds, schooling fish, and every insect society are examples of emergent behaviors - for an excellent review see [7]. Nowhere in the animal’s DNA are these group behaviors coded, nowhere in their lifecycles are these behaviors learned. They arise (emerge) naturally from the interactions of the individuals. For instance, schooling fish can be described as the interaction of individuals, each having these rules:

1. Head toward center of mass
2. Align vectors
3. Don’t crash into your neighbors

Another example is formation flight:

1. Follow leader
2. Fly off either wing of leader
3. Maintain altitude
4. Don't crash into your neighbors

From this the magnificent "V" formations of geese (and aircraft) arise. Yes, human formation flight can be looked at as an emergent behavior (and so can free-market economies [10, Note 2]. Simply put, an emergent behavior is the observed result of the interactions of many individuals making up the society that one might not predict. Emergent behaviors are implicit, no overt plans run them, they just are. This is in contrast to a deliberate behavior where plans are developed, and efforts coordinated amongst individuals a priori to ensure a specific result. Deliberate behaviors are explicit, they are because we want them so.

Emergent behaviors have evolved over millions of years since they assisted the survival of the species. They allow fairly stupid individuals to exhibit complicated group behavior (or at least exhibit it to us since they don't know they are exhibiting this behavior - they don't even realize they exist) benefiting the society in whole. These complicated behaviors are well beyond the capability of any one individual to accomplish (for instance ants only have about a million neurons in their bodies). No plans are required; no learning has to take place, just the proper mix of interacting reactive individual behaviors carried in the individuals DNA [Note 9].

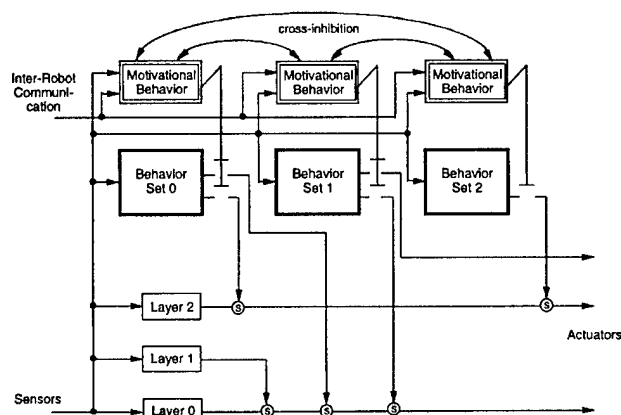
Emergent behaviors have very attractive attributes for inclusion into autonomous UAVs:

- *Inherently Decentralized*: It eliminates the need for any centralized command, an attribute that is desired for some of our distributed UAV control system development.
- *Inherently Implicit*: It eliminates any requirement to explicitly control the individuals. Control is implicit, saving code size that saves money.
- *Inherently Resilient*: Natural history shows that it has ensured the survival of millions of species over millions of years. It is tolerant to imperfection and variance amongst individuals.
- *Inherently Scalable*: Emergent systems are very robust to the addition and subtraction of members, so such systems could work well in high risk situations.

As history shows, simple autonomous systems have successfully used emergent behavior over millions of years to ensure their survival and proliferation. One could arrive at the conclusion that if one was designing a distributed system of simple autonomous entities one should look to emergent behavior as the autonomous control method of choice. Proven in the field, robust to change.

Of course, emergent behavior has its drawbacks:

- *Results are probabilistic*. One cannot say exactly when an outcome will happen, just that there is a probability distribution that it will occur within some interval which depends on the number of individuals, their micro behaviors and the environment.
- *Execution is non-deterministic*. One cannot know exactly how each individual will do something, or how the particular individual's contribution will contribute to the whole, just that the collective will emerge the behavior "somehow".
- *Chaotic behaviors possible*. Depending on the micro behaviors, results could be chaotic (small changes in input result in huge differences in output – the difference between a bee swarm gathering pollen and attacking you).
- *Design is problematic* – non-linear (possibly chaotic) and non-deterministic interacting behaviors relying on environmental aspects cannot be predicted, only observed. Design of such systems is by trial and error (evolution).



Judicial Behavior-Based Architecture

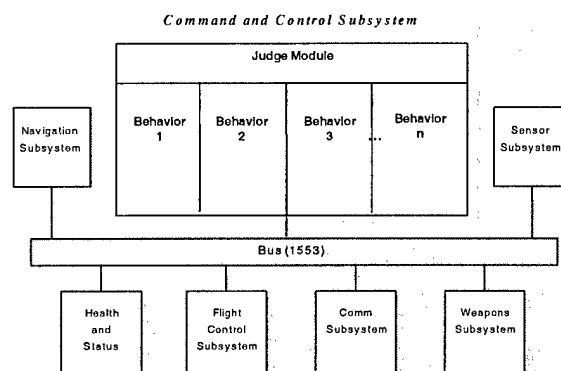


Figure 2: Subsumptive Architecture From Oak Ridge National Laboratory And Judicial Behavioral Control Architecture

- *Design is embedded in the environment.* The fitness of the system cannot be separated from the environment in which it operates.

All of the above conspire to make an emergent system extremely tough to a priori design. Emergent systems we've designed to date have used sets of ad hoc and heuristic rules to develop initial solutions that have proven interesting to "debug" [Note 3].

Reactive Individual Behaviors

The emergent behavior examples in the above come from the given reactive micro behaviors. Reactive? Why reactive? Simple – no plans needed. Something happens, we sense it, and then take action based on our sensory inputs. Most natural systems are purely reactive, only some higher mammals have the capability of deliberate actions - the capability to plan - to assist their survival. Most man-made control systems are this way also, simply reacting to external inputs or disturbances to maintain an output variable (or vector) within some range. Even with our great capability to plan most human actions are driven by reactive behaviors (subconscious, automatic, and motor actions of our neural systems are all reactive). Reactive behaviors mean that we have no plans; therefore we don't need to carry along internal models to assess how we are following those plans. Our control systems can be very simple indeed.

Behavior Architectures

The reactive micro behaviors aren't the only thing required. We also need a way to organize these micro behaviors within individuals such that they can automatically switch behavior executed based on sensory inputs. Just as insects can switch from foraging to nest building behaviors our UAVs must be able to switch behaviors to be versatile. Several of these behavior architecture have been around for a while.

One is a subsumptive architecture [8]. In this architecture some behaviors can subsume (override) others. A good example is what happens when one mistakenly touches a hot surface while watching TV. The current behavior, watching CNN, is subsumed by the survival instinct of reacting to pain. Ever tried to open your eyes while you sneeze? You can't because your body overrides conscience efforts to keep them open. Same idea. Figure 2 shows a simple subsumptive architecture. This details a robotic architecture that Lynne Parker is currently evolving at Oak Ridge National Laboratory [9].

Another architecture used in several of our experimental systems relies on an overarching “Judge” to look at external conditions and choose a behavior to execute. All the behaviors are constantly “petitioning” the Judge, but only one gets to direct the system. This type of architecture is also shown in Figure 2.

Simple animals (as well as ourselves) use such architectures to arrange their behavior sets. We aircraft control system designers have also used subsumptive methods to develop flight control systems for years, developing mode switches for the system in response to failures. Reconfigurable systems modify control laws and limit responses (instill new behaviors) based on sensed failures.

Local Sensing With Sensor Variance Among Individuals

Emergent behaviors result from the interactions of individuals running reactive behaviors in their “control “ systems. But what is required to switch behaviors, or sets of behaviors (such as from “school” to “acquire food” in fish). Something has to trigger this, and the trigger is the sensory inputs for each individual. A common trait of most multi-cellular animal, and all higher order animals, is the ability to sense the environment in several ways. All the sensors do is to inform us of the externals in order to change behaviors - to run a different set of rules. This is very (perhaps directly) analogous to when a finite-state machine senses changes in the system being controlled and switches to a different state in response.

No global knowledge is required to accomplish the task, just information gleaned from interacting with the local environment. Emergent behaviors require only local, versus global, information, therefore they do not require long range communication, nor do they require confirmation that the message was received. In fact, many emergent behaviors are accomplished via stigmergent communication – information acquired via interaction with the environment or task. This is how ants do it with pheromones. All an ant is indicating with a pheromone deposit is “I was here”. That’s all, and yet that is used to run the logistics for an ant colony.

One more point, the robustness of emergent behaviors seems to result from the variance of sensor performance between individuals [7]. The sensors aren’t perfect, and the overall society uses that to assist survival. As externals change, not all individuals change micro behaviors at the same level of change. There is a distribution of sensed values that will cause individual behaviors to change. In practice, this leads to a smooth transition between emergent behaviors, allows several emergent behaviors to co-exist at the same time (upon which more complicated, higher, emergent behaviors can develop), and reduces chaotic tendencies and limit cycling.

Autonomous Entity Societies – The Integration Of The Above

The rich animal world we see around us is largely due to interacting individuals with subsumptive reactive behaviors that emerge the macro behaviors we see. These behaviors are robust and “cost effective” (the natural world is a minimum energy system, so the survivors tend to have the most efficient methods to organize efforts). We want to take advantage of these ideas, be bio-inspired to develop autonomous UAV control systems. However, before we launch into throwing resources at the challenges, we think it’s appropriate to lean back and examine what emergent behavior could do. Are there any simplistic rules as to what this is good for, and what not? Where and when is it appropriate to use swarming? To do this we need to examine what natural systems that use emergent behavior are good at, and that leads us back to examining their attributes.



Swarms Vs. Teams

Comparison Of Attributes



Attribute	Swarm	Team
Temporal	Reactive	Predictive
Composition	Homogeneous	Heterogeneous
Interrelationships	Simple	Complex
Predictability	Probabilistic	Deterministic
Individual Worth	Expendable	Critical
Efficiency	Low	High

Figure 3: Comparison of Swarms and Teams

So What Are The Attributes Of A Swarm, And How Does That Differ From A Team?

Team? Where did that come from? Teaming is simply something that deliberate creatures are good at. Teams are deliberate behaviors – each member has a role to accomplish, knows what that is, knows what the other member's roles are, and knows how they relate as the task is accomplished. They have a plan. A team is the antithesis of a swarm. Determining what teams can do well is just as important to using swarming as determining what swarms do well since it bounds the usefulness of swarm behavior. What do you mean by that? Read on. Let's compare attributes we see in naturally occurring teams and swarms.

Comparing System Attributes: Swarm Versus Team

Figure 3 contains a table that compares attributes of swarms and teams found in nature.

- *Temporal*: Swarms are composed of individuals that can only react, not predict. Teams predict. Predicting capabilities require onboard models and plans, reacting doesn't.
- *Composition*: Swarms are composed of homogeneous, or very limited heterogeneous individual types or roles (workers versus soldier ants as an example). Teams are very heterogeneous with distinct roles assigned to distinct individuals.
- *Interrelationships*: Swarms communicate through simple messages with simple meanings, such as an ant's pheromone trail which just means "I was here". Messages are not targeted to any particular individual, but are broadcasted. Teams communicate with higher-level semantic messages, usually directed at a particular individual. Swarms have simple societies with minimal contact between members, while teams exhibit complicated social behaviors.
- *Predictability*: Swarms are probabilistic; their performance can be described by distributions. They are not deterministic. Swarms have no plans. Teams, on the other hand, are deterministic, with a known objective at a known time achieved via a deliberate (deterministic) plan.
- *Individual Worth*: Swarms get their robustness from the similarity of individuals, loss of one will not hamper a team, and so individual worth is low. Teams on the other hand are composed of specialized individuals, loss of one can cripple critical team functions, and so individual worth is high.

- *Efficiency*: In nature, swarms are inefficient compared to ways humans might decide to do the same things. One of the benefits of teaming is that it maximizes efficiency by synergistic use of heterogeneous resources. Swarms, on the other hand duplicate many tasks, and can even work against each other in some instances. Some researchers indicate that ants are productive only 40% of the time [7].

Note this last attribute isn't necessarily a drawback for the ants, since what they lack in efficiency they make up for in robustness. They trade *efficiency* for *survivability*. Their redundancy can survive attrition, where even a loss of a single team member could spell failure. Inefficiency leaves room for improvement in times of need, allows attrition when things go wrong, but it requires more resources.

So what have we proved? In the natural world, simple animals use simple ways of working with each other. They do not know the other swarm members exist (to know you exist you have to have an idea of who/what you are, which means that you have to have built a mental model of your own existence, something that a critter comprised of reactive behaviors can not do); however, trial and error over millions of years has developed the simple behaviors such that they interact to the good of the collective. In fact, if we can make one observation that seems to be pervasive across the animal world in contrasting swarms and teams:

Nature's Rule Of Thumb: Stupid Things Swarm, Smart Things Team

In nature, swarming things are not smart. If they were, they'd team! [Note 4] This indicates to us that swarming is a good way to accomplish complex tasks using multiple simple (cheap) things; in fact, it might be the only way. This also indicates that there is a limit to what emergent behavior can accomplish, and that limit is exemplified by the tasks animals accomplish by emergent behavior. We need to watch nature; it's telling us something.

Time Criticality – Not On The List Above, Why? Glaring Error?

You'll notice the words "time critical" do not appear in the list above. Either swarms or teams can do single time critical tasks. I used the word single since teams are much better than swarms in doing series of tasks that are time sequenced (you have to follow a plan to know what the sequence of events is, right?). But getting back to the single task, one might pick what one would use based on the complexity of the time critical task (i.e., do I have to follow a particular plan), but just the fact that it's time critical isn't enough to make the differentiation.

What one will have to do is determine proper numbers in the swarm. The exact time it takes a swarm to accomplish a task is a probability distribution with the mean growing shorter as numbers are increased [Note 5]. So, what one might do is determine that the task has to be done by time T with a probability of X , so enough swarm members are inserted to give a 99.9% probability of occurrence.

Note that this still doesn't mean that it will be done; just that we think it will happen with a probability of X .

So What Are Swarms Good For?

Recapping, swarm strengths are in

- Executing reactive tasks

- Executing simple tasks
- Using multiple similar agents
- Robustness to attrition

In addition, emergent behavior stability relies on imperfect sensing, which feed into the strengths of inexpensive sensors. Swarms are non-deterministic things, relying on interactions and the implicit to accomplish tasks. Therefore, swarms are not good for deliberate actions; any time plans have to be followed. They are not suited for doing tasks that have absolute relative times between them, or relative task dependencies. Nor are they well suited for accomplishing tasks that are complex enough to require specialists performing explicit roles. Therefore many tasks that we deliberate humans do are not well suited to emergent behavior [Note 6]. War is one of those deliberate tasks, so one would expect that we don't want the results of a war to be emergent. However, at the more local levels of the conflict there are some tasks emergent behavior (swarms) could be useful for. Remember, we want determinism at the top level (win the war). We don't always need determinism at the lowest levels.

This doesn't mean that we can't consider swarms for actions that use specialized teams, just that when we do we have to realize that this will not lead to the best solution. What we may be able to do is reformulate the task – look at it a different way in order to enable emergent behavior to be useful. Maybe we need to adjust our paradigms to fit methods, rather than the methods to fit paradigms?

The following are a few tasks that we feel swarms are well suited for working within our current paradigm set:

Area Search & Attack

Searching areas for something when you don't know exactly where it could be located is one of the “classical” uses of swarms. This has been covered in several papers, some of the best concepts use systems that look natural [17]. This is also the subject of several DOD studies [5].

The random nature of most swarms means that they will cover an area without knowing what it looks like. They might take longer, but a plan is not needed. In fact, random searches are optimal given no a-priori information [5], or that one is given wrong information (rules of information: it is late, it is incomplete, it is, to some part, wrong). In other words, random searches by swarms fill an important part of a commander's toolbox when faced with the fog of war.

Surveillance & Suppression

Suppression is the act of keeping your adversary from doing something by causing them to fear that the moment they do, you'll pounce on them. Counter air defense is a simple example of it. Our enemies learned that turning on air defense radars was an excellent way of attracting anti-radiation missiles, so they left the radars silent much of the time that decreased the threat to our aircrews. We didn't eliminate it - we suppressed it.

Swarms of simple weapons randomly wandering the airspace offer the same sort of suppressive action. The numbers inherent in the swarm provide the surveillance capability, equipping the swarms with various sensors to attack particular targets of interest. The emergent nature of swarming ensures that the movement across the battlespace will be unpredictable by opponents. The enemy will not know when, or if, they would be observed and targeted by a swarm member, thus raising fear, nurturing disorganization and lethargy, and suppressing their actions.

This is not a new idea. There have always been the concepts of releasing small weapons to wander the battlefield looking for targets of opportunity. What makes the situation ripe to exploit is the development of netted communications, small sensors, more efficient energy storage, and a better understanding of the control algorithms to make emergence work.

Psychological Warfare

Swarms. Just the word conjures up visions of bugs enveloping humans screaming and writhing in agony. There is just something about the human psyche that causes it to cringe when the “swarm” word is used. We fear swarms, whether they be yellow jackets, or harmless gnats. We could leverage that fear to assist psychological operations. Swarms could be used to go against irregular units, militia, or hostile crowds to disperse or suppress. We could make sure the swarming vehicles look the part, making them mimic invertebrates humans naturally fear, such as wasps or spiders. We could make them look nasty. Released against humans, especially humans engaged in their own emergent behavior – mobs, they could start another human emergent behavior, a mass panic [Note 8].

Diversion

Let’s face it, when the yellow jackets are swarming around your picnic, you aren’t thinking of eating, you’re concentrating on getting rid of them. You are not really fearful of them, but you want to eat your lunch (conduct operations) without being bothered. We can leverage this natural tendency of humans to use swarms to create diversions for the enemy that will help our operations. As with the physiological reason above for using swarms, this will work best on the general populace versus trained troops, however, even good troops are susceptible to diversions.

Software Reduction

This isn’t about what swarms can do against humans, but what the use of emergent behavior can do to reduce the amount of code we require in control systems. Swarms make use of the implicit, they do not require plans or deliberate actions, and thus we can eliminate models and database requirements. This is especially important since the cost of developing software for flight critical systems is escalating. Figure 4 shows both the code and cost increases over the last twenty years. We have gone from 100K lines of code developed at a rough cost of \$200 per line of higher order code (HOL), to 400K at \$500 per HOL in our latest EMD systems[13]. Current research systems (tested 6.3) are pushing over a 1M HOL, and systems coming out of 6.2 research have possible sizes (all software agents fully populated) at least 1.8M HOL to over 4.5M HOL, and the year 2002 costs are pushing \$600/HOLOC. At the current cost of testing and certifying vehicle management systems developing such a system is cost prohibitive; however, that must be done to insure system safety [11]. Leveraging bio-inspired control system concepts promises to dramatically reduce code that could dramatically reduce software production costs. We have goals of reducing our software development costs 20% in the next few years. Better coding automation as well as automated testing is dropping costs. Using emergent behavior we have the possibility of removing entire blocks of code, a possibility we are actively pursuing.

Survivability Enhancement

As mentioned above, swarms movement is non-deterministic; we can leverage this to make our systems more survivable. Ever try to hit an individual blackbird in a flock? The random motion of swarming UAVs can be used to confuse gunnery solutions that expect linear, or at least curvilinear, determinism. Remember, even if the individual motion is random, the

flock still gets to its destination, so we can still pursue a group goal with determinism. So, swarming can be used to confuse prediction systems of gunnery solutions both by random motion and by presenting too many targets. Schools of fish use this coupled with their motion and numbers to confuse predators. Such random motion will come at the expense of optimal fuel performance since fuel must be expended doing the random motion (remember efficiency?).

There very well might be much more we can use a swarm for. It is quite possible that by examining a task we find deliberate today we can find ways of doing it emergent – an activity we’re pursuing.

Swarm Technology Questions

One doesn’t get something for nothing. As with anything else, trade-offs are required. As stated in the “Technical Requirements” section many paragraphs before, there are a few challenges for implementing swarms, and these must be faced as we press ahead in our research. Some of the more critical questions are:

How Do I Design Them Efficiently?

Swarming is an emergent behavior. Emergent behavior is highly non-linear, and non-deterministic. It does not have a closed-form solution mathematically. Given these traits, how do we design it into a system?

Well, if I wandered into the office today, pulled aside an engineer, and asked he/she how it was done, he/she would say that they take a look at the task, heuristically determine what micro behaviors might lead to such a global behavior, code them into a model, then watch it miserably fail in the simulation. They failed, but they learned and made the changes in such a way that eventually the model and simulate process leads them to a set that performs adequately. All would like a better way. They realize that it will be trial and error, but something less tedious is desired.

Our solution is to do this the same way nature does – using evolution. Evolution, in this case, is evolutionary computation, using the computer to accomplish the trial and error to develop a productive global behavior. This isn’t a new idea, journals of evolutionary robotics and artificial life contain mounds of articles dedicated to finding the most efficient ways of breeding emergent systems. We are leveraging these developments to accomplish risk reduction for our swarming UAV autonomous controls toolkit. We are not there yet, but we know which way to press.

How Do I Prove They Will Do What I Want, And Prove They Won’t Do Anything Stupid?

Swarming is non-deterministic. One cannot get around that. This means that what a swarm (emergent behavior) does is best described by a probability distribution, and one can never say that you know what a swarm will do, just that you know what it does most of the time. This presents a problem using emergent behavior in several areas we’re interested in, such as flight and safety critical applications where one has to prove beyond a reasonable doubt that a system will do what it’s designed to, and do nothing that will harm friendlies.

Are we dead in the water here? No. First of all, we might delegate swarming to “kill-box” situations in which we don’t care what it attacks, but we believe that is too limiting. Our methods and techniques of system verification and validation are predicated on deterministic systems. We need to change these to allow us to say that the system will be reliable within a certain delta. It doesn’t need to be perfect, just good enough. This is revolutionary not because we’re saying system errors are inevitable, but because we’re saying that the software isn’t perfect, and that there is a finite chance that it will do something unplanned and possible injurious to us.



Autonomy Brings VMS Software Size & Cost Explosion




Figure 4: Software Cost And Size Growth For Vehicle Management System Operational Flight Programs Over The Last Twenty Years And Projections For 2007

We are fully committed to changing the way we look at systems, and the way we verify and validate them to take into account non-determinism [14]. Several efforts are underway to understand how we have to change our paradigms and system evaluation techniques and processes to accommodate swarming and other emergent behavior in our autonomous system toolkit.

How Do I Keep Them From Attacking Friendlies?

Most of the swarming agents we're considering here are pretty dumb (that's why we want to use swarming in the first place), so how do we make sure that they only go against the bad guys? One can look at this as possibly a special case of the question above. Possibly the best way is to physically separate swarms from friendlies, but that may not be possible. One might also be able to "tag" the friendlies with an identification method, possibly similar to pheromones, where the swarm members know not to attack, suppress, scare, etc. Another possibility would be to tag the bad guys, make them stand out to be specifically targeted, such as killer bees know to attack the closest thing emitting CO₂. We could ask for operator intervention before taking lethal action; however, if there are a swarm of agents in a target rich environment this might overwhelm operators, not to mention the communication system. All of these, and more, are being investigated as we press forward to insure our swarms are safe, or at least as safe as a lethal system can be.

How Can I Make Them More Efficient, But Retain Robustness?

This question comes up a lot, especially when your background is in optimal control theory. Swarms, by their very nature, are inefficient, but we humans like efficiency, so can we find a way to retain the good characteristics (robustness) while decreasing bad characteristics (inefficiency)?

We believe the answer to this is “no, in general you can’t”. The same things that make a swarm robust make it inefficient. There is a direct trade off between performance optimality (efficiency) and robustness [Note 7]. Robustness is achieved at a price. Rather than trying to make it more efficient, we need to make sure that emergent behaviors are used for suitable applications.

We will, however, keep track of efficiency to generate a track record as well as lessons learned. Doing so will give us an idea of how inefficient swarming is such that we can plan for energy storage and management systems.

What Is The Critical Mass Needed For Swarming To Occur?

One thing can’t swarm, multiple things can. For an emergent behavior to operate individuals must interact, which indicates that there could be critical masses for the different global behaviors of interest. Up to a point, one might expect that the time to accomplish a task might increase with decreasing swarm membership, and this is borne out by experiments (research also points out that for some tasks, there is a maximum swarm membership, above which task efficiency actually drops off – Note 5). Research also shows that below some swarm population number (task dependent) emergent behaviors break down. This critical mass must be determined to insure the proper amounts of members are inserted to make up the swarm, allowing for attrition, dispersal, etc. This will drive the affordability of the overall system. Guess what, this is another area where no theory exists; systems must be designed using trial and error - another place where we look to evolutionary computation to help us come up with answers.

How Do I Change Swarm Behavior Real-Time?

As a human operator, I want to be able to change tasking to my UAVs real time. This is very important due to the unpredictability of hostilities. As a systems designer, I want my systems to be multi-purpose to increase affordability. The offshoot of this is I want to enable swarms to be flexible in their behaviors based on the mission, and if the mission changes, I would like to re-task them.

Assuming that I do not embed higher functionality into the swarm members (such as being able to figure out which emergent behavior to use via a plan), I can either try to figure out all the possible uses and global behaviors a priori and set up the local rules to accommodate them, or I can change them “on the fly”. We are pursuing both methods. The first requires more memory to store the sets of local rules along with the sensing to determine the changed environment, the second mean that the operator has to feed a new set of rules to each swarm member simultaneously, requiring a good comm system on the UAVs. (And what happens if only part of the swarm gets the new rules – what interesting behaviors might emerge from that?) The new emergent behaviors can either be canned, waiting on line for insertion, or we could use evolutionary computation to develop them real-time. Which is best? Don’t know – this is an active research area for us.

Summary

Swarming UAVs, although not a general panacea for all possible conflict situations we might find ourselves in, can fit niche situations and can provide interesting options on how to pursue mission objectives.

We've defined swarming as an emergent behavior, relying on the interactions of individuals running simple local rules. It depends on the local agents having reactive rules arranged in a subsumptive architecture. Swarming's attributes of simplicity, robustness and scalability, as well as its drawbacks of inefficiency and non-determinism, lead us to believe it's well suited for missions involving

- Area search and attack where target distribution and location are not known a priori
- Surveillance, diversion, and suppression of hostile force's actions
- Psychological warfare
- System software complexity reduction

We are actively searching the applicability region for emergent behaviors in UAV autonomous control. We are especially attracted to the software cost reductions proper application of this technology could realize. Proving non-deterministic systems are safe and effective to realize this affordability increase is a challenge that we are tackling over the next several years as we develop swarms.

Swarms – we used to call it the “S” word, another buzzword to put on management charts, but our research is making it a much more useful technique in autonomous control. Soon the only “buzz” about it will be the sound of our vehicles accomplishing their missions successfully.

Acknowledgments

I would like to thank the AFRL Control Automation team for their review, critiques, and insight while developing this paper.

References

1. Hitchcock, A. *The Birds*. 1963
2. *Them*. Warner Brothers. 1954
3. *Damnation Alley*. 20th Century Fox Studios. 1977
4. Kinsler, R. *Taxonomy of Animal Behaviors Relating to Unmanned Systems*. White paper being readied for publication as AFRL Technical Report. August 2001.
5. Frelinger, Kvitky, Stanley. *Proliferated Autonomous Weapons – An Example of Cooperative Behavior*. Rand Corp. briefing to USAF. May 1997. Available on the web from: <http://www.rand.org/publications/DB/DB239/DB239.pdf>
6. Merriam-Webster's Dictionary. Available on line from <http://www.m-w.com/home.htm> .
7. Bonabeau, E., et al. *Swarm Intelligence: From Natural to Artificial Systems*. Santa Fe Institute. Oxford University Press. July 1999.
8. Brooks, R., Flynn, A. *Fast, Cheap And Out Of Control: A Robot invasion Of The Solar System*. Journal of The British Interplanetary Society, Vol. 42, pp 478-485. 1989.
9. Parker, L. *ALLIANCE: An architecture for fault-tolerant multi-robot cooperation*. IEEE Transactions on Robotics and Automation, 14 (2):220-240. 1998.
10. Waldrop, M. *Complexity – The Emerging Science at the Edge of Order and Chaos*. Touchstone Book published by Simon & Schuster, New York. 1992.

11. Clough, B. *Autonomous UAV Control System Safety – What Should It Be, How Do We Reach It, And What Should We Call It?* NAECON 2000 paper. October 2000
12. Clough, B. *Metrics, Schmetrics! How the heck can one figure out how autonomous a UAV is?* Presentation at the 26th Dayton-Cincinnati Aerospace Symposium and unpublished paper. March 2001
13. Clough, B. *Unmanned Aerial Vehicles: Autonomous Control Challenges, A Researcher's Perspective*. World Aviation Congress 2000, San Diego, CA October 2000
14. Clough, B. *Autonomous UAVs – Give Them Free Will And They Will Execute It!* Proceedings of the AUVSI Unmanned Systems 2001 Symposium. Baltimore, MD. 2001.
15. Goldberg, P. "Variability in Brain Function and Behavior", published in *The Encyclopedia of Human Behavior, Volume 4*, Academic Press, 1994 (pp 447-458)
16. Arquilla, Ronfelt. *Swarming & The Future of Conflict*. Rand Corp. 2000. Available on the web from <http://www.rand.org/publications/DB/DB311/DB311.pdf>
17. Tilden, M. *Coordinated Covert Detector Ecologies*. Los Alamos National Laboratory. Submitted to ASME Symposium on Micro Robot Design and Control. November 1997.

Notes:

1. I managed to get clear before the firecracker went off, but wouldn't come close until later that night. The hornets were a bit agitated.
2. In fact, one can argue that the eventual collapse of all purely socialist societies is due to socialism not being an emergent behavior of interacting intelligent selfish beings (humans) while capitalism is! Something Karl Marx never thought about...
3. During the course of developing autonomous UAV formation flight agents (software) we decided to use emergent behavior to reduce code size and global information sharing requirements. Initial simulations showed behaviors far different than what we thought we were going to get. The designer thought that he had captured all the requirements and nuances in the local rules. Turns out that several other factors that nobody would have predicted beforehand impacted the final software design. We had to go back and use trial and error to eliminate these. An interesting confirmation of the difficulties building desired emergent behaviors.
4. This probably all gets back to entropy – in some Gaia kind of way the universe has so much invested in us, and we are not that many, that we have to team to survive. Our intelligence leads to teaming which leads to survival, which leads to greater intelligence, and so forth. No wonder most emergent behaviors of humans are unproductive.
5. This holds up to a point. Depending on the task, getting to many members of a swarm could hamper operations. There is an optimal range of swarm members for particular types of tasks that can only be determined by trial and error.
6. However, such things as riots and bar room brawls are emergent behaviors of humans, as well as the stock market and urban sprawl, so maybe we aren't that deliberate after all.
7. Oh, I suppose one could argue that you can have robust optimal techniques, especially since we've spent so much effort over the last twenty years developing them (I'm from a control theoretic background so I can see the questions coming). The fact is if the system is optimized, then it can't be robust almost by definition. This is especially true of linear systems. For instance, there is a direct relationship between the quality factor "Q" of a filter and the bandwidth it will work over. High Q filters are very optimal at eliminating a frequency, but they are intolerant of frequency deviations on the signal to filter. Low Q filters don't filter as much, but they work over wider bandwidths. Another way of looking at this (inspired by an event while writing a different paper) is that an extension cord optimized to reduce the length of wire from outlet to laptop PC is not robust to a 3-year old daughter tripping over it, but one with 6 ft of extra cord is! Robustness requires us to back off on performance.
8. I was careful to point out this was to go against untrained humans. This probably won't work against well-trained troops since part of the training is to eliminate fear. Against these types of targets swarms could be used for another purpose, diversion...
9. One could really say that the reactive behaviors are not coded in the DNA, just that the DNA is set up such that the behaviors self-organize, that the individual behaviors are an emergent phenomena of the interactions of the cells in each individual. This means that emergent phenomena on one scale are dependent on emergent phenomena on other scales, or that in a way emergence is scale independent.